

Enabling Dark Energy Science with Deep Generative Models of Galaxy Images

Siamak Ravanbakhsh¹, François Lanusse², Rachel Mandelbaum², Jeff Schneider¹, and Barnabás Póczos¹

¹*School of Computer Science, Carnegie Mellon University*

²*McWilliams Center for Cosmology, Carnegie Mellon University*

Abstract—Understanding the nature of dark energy, the mysterious force driving the accelerated expansion of the Universe, is a major challenge of modern cosmology. The next generation of cosmological surveys, specifically designed to address this issue, rely on accurate measurements of the apparent shapes of distant galaxies. However, shape measurement methods suffer from various unavoidable biases and therefore will rely on a precise calibration to meet the accuracy requirements of the science analysis. This calibration process remains an open challenge as it requires large sets of high quality galaxy images. To this end, we study the application of deep conditional generative models in generating realistic galaxy images. In particular we consider variations on conditional variational autoencoder and introduce a new adversarial objective for training of conditional generative networks. Our results suggest a reliable alternative to the acquisition of expensive high quality observations for generating the calibration data needed by the next generation of cosmological surveys.

The last two decades have greatly clarified the contents of the Universe, while leaving several large mysteries in our cosmological model. We now have compelling evidence that the expansion rate of the Universe is accelerating, suggesting that the vast majority of the total energy content of the Universe is the so-called *dark energy*. Yet we lack an understanding of what dark energy actually is, which provides one of the main motivations behind the next generation of cosmological surveys such as LSST (LSST Science Collaboration et al., 2009), Euclid (Laureijs et al., 2011) and WFIRST (Green et al., 2012). These billion dollar projects are specifically designed to shed light on the nature of dark energy by probing the Universe through the *weak gravitational lensing* effect –i.e., the minute deflection of the light from distant objects by the intervening massive large scale structures of the Universe. On cosmological scales, this lensing effect causes very small but coherent deformations of background galaxy images, which appear slightly *sheared*, providing a way to statistically map the matter distribution in the Universe. To measure the lensing signal, future surveys will image and measure the shapes of billions of galaxies, significantly driving down statistical errors compared to the current generation of surveys, to the level where dark energy models may become distinguishable.

However, the quality of this analysis hinges on the accuracy of the shape measurement algorithms tasked with estimating the *ellipticities* of the galaxies in the survey. This point is particularly crucial to the success of these missions, as any unaccounted for measurement biases in their ensemble

averages would impact the final cosmological analysis and potentially lead to false conclusions. In order to detect and/or calibrate any such biases, future surveys will heavily rely on image simulations, closely mimicking real observations but with a known ground truth lensing signal.




Fig. 1: Illustration of the processes involved in the measurement of weak gravitational lensing. The light from distant galaxies is deflected by the matter in the Universe, causing a shearing of the galaxy images, which are then further blurred by the atmosphere and the telescope optics and finally pixelated into a noisy image by the imaging sensor. Image credit: Mandelbaum et al. (2014), adapted from Kitching et al. (2010).

Producing these image simulations, however, is challenging in itself as they require high quality galaxy images as the input of the simulation pipeline. Such observations can only be obtained by extremely expensive space-based imaging surveys, which will remain a scarce resource for the foreseeable future. The largest current survey being used for image simulation purposes is the COSMOS survey (Scoville et al., 2007), carried out using the *Hubble Space Telescope* (HST). Despite being the largest available dataset, COSMOS is relatively small, and there is great interest in increasing the size of our galaxy image samples to improve the quality of this crucial calibration process.

In this work, we propose an alternative to the expensive acquisition of more high quality calibration data using deep conditional generative models. In recent years, these models have achieved remarkable success in modeling complex high-dimensional distributions, producing natural images that can pass the visual Turing test. Two prominent approaches for training these models are **variational autoencoder** (VAE) (Kingma and Welling, 2013; Rezende et al., 2014) and **generative adversarial network** (GAN) (Goodfellow et al., 2014). Our aim is to train a conditional variation of these models using existing HST data and generate new galaxy images




Fig. 2: Samples from the GALAXY-ZOO dataset and generated samples using conditional generative adversarial network of Section III. Each synthetic image is a 128×128 colored image (here inverted) produced by conditioning on the same set of features $y \in [0, 1]^{37}$ as its real pair. These instances are selected from the test-set and were unavailable to the model during the training.

conditioned on statistics of interest such as the brightness or size of the galaxy. This will allow us to synthesize calibration datasets for specific galaxy populations, with objects exhibiting realistic morphologies.

In the following, Section I gives a brief background on the image generation for calibration and its significance for modern cosmology. We then review the current approaches to deep conditional generative models and introduce new techniques for our problem setting in Sections II and III. In Section IV we assess the quality of the generated images by comparing the conditional distributions of shape and morphology parameters between simulated and real galaxies, and find good agreement.

I. WEAK GRAVITATIONAL LENSING

In the weak regime of gravitational lensing, the distortion of background galaxy images can be modeled by an anisotropic shear, noted γ , whose amplitude and orientation depend on the matter distribution between the observer and these distant galaxies. This shear affects in particular the apparent ellipticity of galaxies, denoted e . Measuring this weak lensing effect is made possible under the assumption that background galaxies are randomly oriented, so that the ensemble average of the shapes would average to zero in the absence of lensing. Their apparent ellipticity e can then be used as a noisy but unbiased estimator of the shear field γ : $\mathbb{E}[e] = \gamma$. The cosmological analysis then involves computing auto- and cross-correlations of the measured ellipticities for galaxies at different distances. These correlation functions are compared to theoretical predictions in order to constrain cosmological models and shed light on the nature of dark energy.

However, measuring galaxy ellipticities such that their ensemble average (used for the cosmological analysis) is unbiased is an extremely challenging task. Fig. 1 illustrates the main steps involved in the acquisition of the science images. The weakly sheared galaxy images undergo additional distortions (essentially blurring) as they go through the atmosphere and telescope optics, before being acquired by the imaging sensor which pixelates the noisy image. As this figure illustrates, the cosmological shear is clearly a subdominant effect in the final image and needs to be disentangled from subsequent blurring by the atmosphere and telescope options. This blurring, or Point Spread Function (PSF), can be directly

measured by using stars as point sources, as shown at the top of Fig. 1.

Once the image is acquired, shape measurement algorithms are used to estimate the ellipticity of the galaxy while correcting for the PSF. However, despite the best efforts of the weak lensing community for nearly two decades, all current state-of-the-art shape measurement algorithms are still susceptible to biases in the inferred shears. These measurement biases are commonly modeled in terms of additive and multiplicative bias parameters c and m defined as:

$$\mathbb{E}[e] = (1 + m) \gamma + c \quad (1)$$

where γ is the true shear. Depending on the shape measurement method being used, m and c can depend on factors such as the PSF size/shape, the level of noise in the images or, more generally, intrinsic properties of the galaxy population (like their size and ellipticity distributions, *etc.*). Calibration of these biases can be achieved using image simulations, closely mimicking real observations for a given survey but using galaxy images distorted with a known shear, thus allowing the measurement of the bias parameters in Eq. (1).

Image simulation pipelines, such as the *GalSim* package (Rowe et al., 2015), use a forward modeling of the observations, reproducing all the steps of the image acquisition process in Fig. 1, and therefore require as a starting point galaxy images with high resolution and S/N. The main difficulty in these image simulations is therefore the need for a calibration sample of high quality galaxy images representative of the galaxy population of the survey being simulated. Our aim in this work is to train a deep generative model which can be used to cheaply synthesize such data sets for specific galaxy populations, by conditioning the samples on measurable quantities.

A. Data set

As our main dataset, we use the COSMOS survey to build a training and validation set of galaxy images and extract from the corresponding catalog a condition vector y with three features: half-light radius (measure of size), magnitude (measure of brightness) and redshift (cosmological measure of distance). To facilitate the training, we align all galaxies along




Fig. 3: Samples from the COSMOS dataset and generated samples using the conditional variational autoencoder (C-VAE, scheme I) and our variation on conditional generative adversarial network (C-GAN). Each column image shows three 64×64 images (here inverted) produced by conditioning on the same set of features $y \in \mathbb{R}^3$ in the test-set. Due to its high dynamic range, most figures are very faint. In the bottom three rows, each image is individually normalized.

their major axis and produce 85,000 instances of 64×64 image stamps using the GalSim package.

We also use the GALAXY-ZOO dataset (Willett et al., 2013) to demonstrate the abilities of our alternative conditional adversarial objective. Each of the 61,000 galaxy images in this dataset is accompanied by $y \in [0, 1]^{37}$ features produced using a crowd-sourced set of questions that form a decision tree. We cropped the central 50% of these images and resized them to 128×128 pixels. We augmented both datasets by flipping the images along the vertical and horizontal axes.

II. CONDITIONAL VARIATIONAL AUTOENCODER

Applications in semi-supervised learning and structured prediction have motivated different versions of the “conditional” variational autoencoder (C-VAE) in the past (Kingma et al., 2014; Sohn et al., 2015). Although the objective and architecture that we discuss here is similar to those of (Kingma et al., 2014; Sohn et al., 2015), there are some differences due to different objectives.

We are interested in learning the conditional density $p^*(x | y)$ for $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, given a set of observations $\mathcal{D} = (\hat{x}^1, \hat{y}^1), \dots, (\hat{x}^N, \hat{y}^N)$, by learning θ that maximizes the conditional likelihood $\prod_{(\hat{x}, \hat{y}) \in \mathcal{D}} p_\theta(\hat{x} | \hat{y})$ – e.g., for the COSMOS dataset $\mathcal{X} = \mathbb{R}^{64 \times 64}$ and $\mathcal{Y} = \mathbb{R}^3$. In a *latent-variable model*, an auxiliary variable $z \in \mathcal{Z}$ is introduced to increase the expressive power of $p_\theta(x, z | y)$, such that $\int_{\mathcal{Z}} p_\theta(x, z | y) dz$ is the marginal of interest. Here, different assignments to z can explain variations and complex statistical dependencies in $p(x | y)$.

To enable efficient (ancestral) sampling from this model, p_θ can be a directed model $p_\theta(x, z | y) = p_{\theta_1}(z | y) p_{\theta_2}(x | z, y)$, where we first sample $z \sim p_{\theta_1}(\cdot | y)$ followed by $x \sim p_{\theta_2}(\cdot | z, y)$. An expressive form for the conditional distributions p_{θ_1} and p_{θ_2} is a deep neural network, that can represent complex directed graphical models. Here, for example, we model $p_{\theta_2}(x | z, y)$ using multi-layered convolutional or densely connected neural networks that encode the mean and variance of a multi-variate Gaussian for the COSMOS dataset and the expectation of Bernoulli variables for the GALAXY-ZOO dataset.

To learn the parameters θ one needs to estimate the posterior $p_\theta(z | x, y)$, which is often intractable in directed models. An elegant solution to this problem is to introduce a second




Fig. 4: Cross-correlation between y and z in C-VAE when $p(z | y) = p(z)$, with and without cross-correlation penalty.

directed model $q(z | x, y)$, called inference or *recognition model*. This conditional distribution is also encoded as a deep neural network and it is tasked with estimating the intractable posterior $p_\theta(z | x, y)$.

This is achieved through a **variational bound** on the conditional log-likelihood:

$$\log(p_\theta(\hat{x} | \hat{y})) \geq -\mathbb{D}_{\text{KL}}(q_\phi(z | \hat{x}, \hat{y}) \| p_{\theta_1}(z | \hat{y})) + \mathbb{E}_{z \sim q_\phi(\cdot | \hat{x}, \hat{y})} [\log p_{\theta_2}(\hat{x} | z, \hat{y})] \quad (2)$$

where the first term is the KL-divergence between the posterior q_ϕ and the conditional prior p_{θ_1} and the second term is the reconstruction error – that is we want the model to achieve low reconstruction error while encoding the dataset. At the same time the KL-divergence term encourages the code to follow a distribution, dictated by the condition y . Fortunately, reparametrization-trick (Kingma and Welling, 2013; Rezende et al., 2014; Williams, 1992) enables maximization of this lower-bound (i.e., learn θ_1 , θ_2 and ϕ) using stochastic back-propagation through the layers of these three neural networks. This enables maximizing the log-likelihood of an expressive model with large number of parameters through variations of stochastic gradient descent.

A. Cross-Correlation

Inspired by application of cross-correlation in disentangling the factors in a autoencoder by Cheung et al. (2014), we also consider an alternative method of conditioning in VAE. Let us proceed with a simple question: what happens here if we simplify the prior $p_{\theta_1}(z | y) \Rightarrow p_{\theta_1}(z)$? In principle, the simplified C-VAE would try to make the posterior $q_\phi(z | x, y)$

independent of y .¹ In this case, for generating samples $x \sim p_\theta(\cdot | y)$, we could still sample $z \sim p_{\theta_1}(\cdot)$ and then generate $x \sim p_{\theta_2}(\cdot, | y, z)$.

In practice, we observe z and y become more and more decorrelated during the training, but this happens at a slow pace. We can further enforce this decorrelation using a mini-batch cross-correlation penalty

$$C(\{\hat{y}\}, \{z\}) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{i,j} \left(\frac{1}{N} \sum_{n=1}^N (\hat{y}_i^{(n)} - \bar{y}_i)(z_i^{(n)} - \bar{z}_i) \right)^2$$

where $\{\hat{y}\}/\{z\}$ are conditions/codes in a *mini-batch* of size N , where $z \sim q_\psi(\cdot | \hat{x})$ and i, j index dimensions of \hat{y}, z respectively. Here \bar{y}_i and \bar{z}_i are mini-batch average values.




Fig. 5: Negative log-likelihood of different C-VAE schemes. Note that scheme II can only serve as a baseline and due to correlation between \hat{y} and z cannot be used for conditional sampling.

Lack of cross-correlation only entails independence, if both y_i and z_i have Gaussian distribution. Although $p_{\theta_1}(z)$ is by design a standard Gaussian, the condition y may have an arbitrary distribution. To resolve this, we transform $\hat{y}_i \rightarrow F_{\mathcal{N}}^{-1}(F_{y_i}(\hat{y}_i))$, where F_{y_i} is the empirical cumulative distribution function (CDF) for $\hat{y} \in \mathcal{D}$ and $F_{\mathcal{N}}^{-1}$ is the (numerically approximated) inverse CDF of Gaussian. The transformed variable has a Gaussian distribution.

B. Experiments

Figure 4 compares the reduction in the average cross-correlation between \hat{y} and z for the same network, with and without the cross-correlation penalty. For numerical stability we linearly increase the penalty coefficient from 0 to 1000 over iterations. These results are for the COSMOS dataset. All C-VAE results are using the log-pixel-intensity, also for numerical stability.

Figure 5 compares $-\log(p_\theta(\hat{x} | \hat{y}))$ for three models:

- I using a neural network to encode $p_{\theta_1}(z | y)$
- II using $p_{\theta_1}(z | y) = p_{\theta_1}(z)$
- III $p_{\theta_1}(z | y) = p_{\theta_1}(z)$ plus cross-correlation penalty

The figure suggests that the first scheme eventually produces better models. It also shows that enforcing the independence of z and y only slightly decreases the likelihood, compared to the baseline II where z and y remain highly dependent.

¹This is because the information content of y is already available to the generative model $p_\theta(x | y, z)$ for reconstruction and reducing the information exchange through z should reduce the KL-divergence penalty $\mathbb{D}_{\text{KL}}(q_\psi(z | \hat{x}, \hat{y}) || p(z))$.

III. A NEW OBJECTIVE FOR ADVERSARIAL TRAINING

A major problem with VAE-generated images is their blurriness. A few recent works address this issue (Kingma et al., 2016; Larsen et al., 2015; Dosovitskiy and Brox, 2016) – e.g., by defining a more expressive reconstruction loss. Fortunately, the noise model is available for COSMOS images, and the added noise to some extent reduces this problem in our application (see Section IV).




Fig. 6: The prediction error for real and generated images in C-GAN for COSMOS dataset.

An alternative to generative modeling that does not suffer from this problem is offered by adversarial training of generative networks (Goodfellow et al., 2014). In the **adversarial setting**, a generator $G_\omega : \mathcal{Z} \rightarrow \mathcal{X}$ attempts to fool the discriminator $D_\psi : \mathcal{X} \rightarrow [0, 1]$ into classifying its fake instances $x = G(z)$ as real, while the discriminator’s objective is to correctly classify the two sources of real versus generated instances. Deep networks representing these adversaries are trained alternatively, and under some conditions p_G (the implicit distribution of the generator G_ω for $z \sim \mathcal{U}(0, 1)$) converges to p^* –i.e., at this fixed-point, the generator produces realistic images that are indistinguishable by the discriminator.

The **conditional variation** of this method was first introduced by Mirza and Osindero (2014) and used in a cascade of conditional models with increasing resolution in (Denton et al., 2015). In these conditional models, the generator $G_\omega : \mathcal{Z} \times \mathcal{Y} \rightarrow \mathcal{X}$ and the discriminator $D_\psi : \mathcal{X} \times \mathcal{Y} \rightarrow [0, 1]$, are both deep neural networks that are now conditioned on the same observed variable $\hat{y} \in \mathcal{D}$. The min-max formulation of this adversarial setting seeks a saddle-point for

$$\min_{\omega} \max_{\psi} \mathbb{E}_{\hat{x}, \hat{y} \in \mathcal{D}, z \in \mathcal{U}} [\log(D_\psi(\hat{x}, \hat{y})) + \log(1 - D_\psi(G_\omega(z, \hat{y}), \hat{y}))]$$

In practice it is much more efficient to use a different loss function for the generator as it produces stronger gradients for the generator at the beginning (Goodfellow et al., 2014):

$$\max_{\psi} \mathbb{E}_{\hat{x}, \hat{y} \in \mathcal{D}, z \in \mathcal{U}} [\log(D_\psi(\hat{x}, \hat{y})) + \log(1 - D_\psi(G_\omega(z, \hat{y}), \hat{y}))]$$

$$\max_{\omega} \mathbb{E}_{z \in \mathcal{U}} [\log(D_\psi(G_\omega(z, \hat{y}), \hat{y}))]$$

Here, one must carefully adjust the expressive power of G and D to avoid oscillations, and domination of either adversary. The choice of hyper-parameters is known to be a




Fig. 7: Comparison of a C-VAE sample before and after adding noise and a real COSMOS image with corresponding size, magnitude and redshift.

major hurdle in training of adversarial networks and using this scheme, despite much effort, we could not train a generator for our problem that uses *continuous* conditional variables.

We introduce an **alternative adversarial objective** for conditional generative modeling that in our experience is more stable and did not require any hyper-parameter tuning in our application. The basic idea is simple: A **predictor** $R : \mathcal{X} \rightarrow \mathcal{Y}$ replaces the discriminator $D : \mathcal{X} \times \mathcal{Y} \rightarrow [0, 1]$. The predictor attempts to produce predictions of the condition $\hat{y} \in \mathcal{D}$ for the real data, that are at least as good as its predictions for generated instances. The generator’s objective is to produce instances with low prediction error

$$\text{Predictor: } \min_{\psi} \min\{0, \quad (3)$$

$$\mathbb{E}_{\hat{x}, \hat{y} \in \mathcal{D}, z \in \mathcal{U}} [\ell(R_{\psi}(G_{\omega}(z, \hat{y}), \hat{y})) - \ell(R_{\psi}(x), \hat{y})] \}$$

$$\text{Generator: } \min_{\omega} \mathbb{E}_{\hat{y} \in \mathcal{D}, z \in \mathcal{U}} [\ell(R_{\psi}(G_{\omega}(z, \hat{y}), \hat{y}))] \quad (4)$$

where in our application $\ell(y, \hat{y}) = \|y - \hat{y}\|_2^2$.

Why should the generator produces realistic images at all as long as the predictor makes equally bad predictions for both real and generated images? Both errors Eqs. (3) and (4) will be low in this case. The key here is that the generator always seeks to improve its samples to increase their prediction accuracy and therefore the dynamics of this adversarial setting does not allow this mode of failure.

This scheme, also relaxes the constraint on the **expressive power of the adversaries**. This is because the predictor has no incentive to lower the error for the real data, as long as its prediction errors are not worse than those of the generated data. Therefore, it is only the generator that fuels the competition and training is practically finished when the generator is unable to improve.

A mode of failure that our scheme does not resolve is the **collapse of generator**, where generator $G(y, z)$ repeats few output patterns by solely relying on y and basically ignoring the random feed z . The predictor eventually realizes this repeating pattern in generated data but gradient descent can no longer rescue the generator from this local optima. A solution to this problem called mini-batch discrimination was recently proposed by Salimans et al. (2016), where each instance in the mini-batch is augmented with information about its differences with other instances in the same mini-batch. Predictor can therefore detect this tendency of the generator early on, and the generator incurs a loss for its behavior before its complete

collapse. For better mini-batch statistics, we use relatively larger mini-batches with 128/256 instances.

A. Experiments

Following Radford et al. (2015) we use (de)convolutional layers with (fractional) stride, batch normalization (Ioffe and Szegedy, 2015) and leaky-ReLU activation functions in our deep networks. For optimization, we use Adam (Kingma and Ba, 2014) with reduced exponential decay rate of .5 for the first moment estimates.

Figure 6 reports the prediction loss $\ell(R_{\psi}(G_{\omega}(z, \hat{y}), \hat{y}))$ and $\ell(R_{\psi}(x), \hat{y})$ for the COSMOS dataset, where we use 4 (de)convolution layers. The figure suggests that the predictor tends to keep the prediction error of the real images slightly higher than that of generated images. Both of these quantities reduce over time, and their agreement with validation errors could monitor convergence. The fact that the error is decreasing over time and prediction error for both real and generated data remains close to each other is due to having a “laid back” predictor – i.e., by removing $\min(0, \cdot)$ operation in predictor’s loss, we would lose both of these properties.

For illustration purposes, we applied the same method to the GALAXY-ZOO dataset. Figure 2 shows some instances in the test-set accompanied by C-GAN generated image conditioned on the same \hat{y} . For this dataset we used 5-layer fully (de)convolutional generator and predictor, mini-batch discrimination, batch-normalization and tanh activation function for the final layer of the generator.

IV. VALIDATION

In this section, we assess the quality of the model generated galaxy images by comparing common image statistics used in weak lensing analyses. Our aim is to consistently measure the same statistics on real COSMOS images and images generated by our model for the same set of input variables y . These statistics are affected by the presence of noise in the image, but as was noted in the previous section, our C-VAE generates essentially noiseless images, which prevents direct comparison with real images. We limit this analysis to C-VAE generated images (as we found it to produce more consistent results compared to C-GAN) and add a noise field to our generated images. This noise model, calibrated for COSMOS observations, is provided by the GalSim package; see Fig. 7.

The most commonly used image statistics in weak lensing analyses rely on the second moments of the galaxy’s intensity profile $I(u_1, u_2)$, where (u_1, u_2) are pixel coordinates. The second moment tensor Q is defined as:

$$Q_{\alpha\beta} = \frac{\int du_1 du_2 W(u_1, u_2) I(u_1, u_2) u_{\alpha} u_{\beta}}{\int du_1 du_2 W(u_1, u_2) I(u_1, u_2)},$$

with $(\alpha, \beta) \in \{1, 2\}$ and where W is a weighting function. This tensor can be used to define a size measurement $\sigma = |\det(Q)|^{1/4}$ which reduces to the standard deviation if the light profile is a Gaussian. More importantly, the second moments are commonly used to measure galaxy ellipticities which can be defined as:

$$e = e_1 + ie_2 = \frac{Q_{11} - Q_{22} - 2iQ_{12}}{Q_{11} + Q_{22} + 2(Q_{11}Q_{22} - Q_{12}^2)^{1/2}}$$

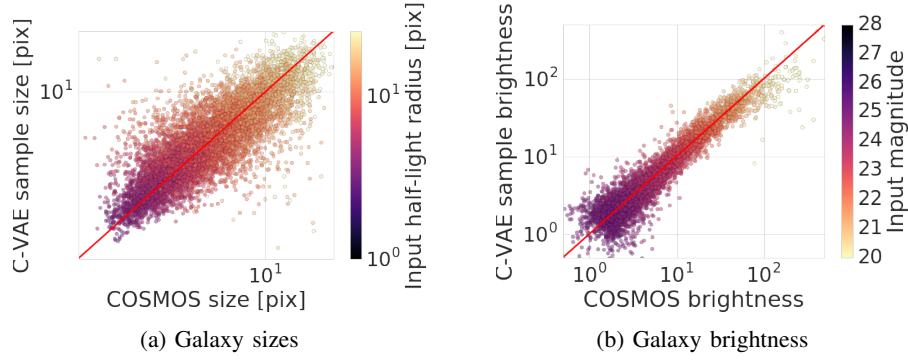


Fig. 8: Comparison of galaxy sizes and brightness between real COSMOS images and C-VAE samples. Colors indicate the value of the relevant variable used to condition the generated images (half-light radius for size and magnitude for brightness)

To measure Q in practice, we use the *adaptive moments* method (Hirata and Seljak, 2003; Mandelbaum et al., 2005) which estimates the second order moments by fitting an elliptical Gaussian profile to the galaxy light profile. As a side product of this method, we can also use the amplitude of the best fit Gaussian model as a proxy for the brightness of the galaxy.




Fig. 9: Comparison of galaxy ellipticity (left) and size (right) distributions measured from second moments between real COSMOS images and CVAE samples.

We compare real COSMOS images to C-VAE samples by processing the images in pairs, where every COSMOS galaxy in our validation set is associated to a C-VAE sample conditioned on the half-light radius, magnitude and redshift of the real galaxy. Fig. 8a shows for each pair of images the galaxy size σ , as measured using second order moments; see also Fig. 3. The color of the points indicates the half-light radius of the COSMOS galaxy in the pair, also used to condition the C-VAE sample. As can be seen, the sizes of generated galaxies are generally unbiased. Fig. 8b shows the similar results for brightness; C-VAE is generating samples of the correct brightness without any significant bias.

The most relevant image statistics for weak lensing science are the ellipticity and size distributions of a given galaxy sample. Fig. 9 compares these overall distributions measured on real and generated galaxies. Note that contrary to the previous test where the quantities considered (size and brightness) were part of the condition variable y , the ellipticity is not. Therefore, this test allows us to check how well the model

is able to blindly learn correct galaxy shapes. This figure shows that despite being slightly more elliptical than real galaxies, the ellipticity distribution of the C-VAE samples is broadly consistent with the COSMOS distribution. Fig. 9 also compares size distributions which are in good agreement. This comes as no surprise however as C-VAE samples are explicitly conditioned on galaxy sizes and the previous test has shown these samples to be largely unbiased.

CONCLUSION

In this paper, we proposed novel techniques and studied the application of two most promising methods for deep conditional generative modeling in producing galaxy images. In the future, we plan to measure more subtle morphological statistics in generated images and find ways for simultaneous learning of the noise model. We are also investigating the application of our variation on adversarial training in other settings and assessing the effectiveness of the predictor as a stand-alone classification/regression model.

REFERENCES

- Brian Cheung, Jesse A Livezey, Arjun K Bansal, and Bruno A Olshausen. Discovering hidden factors of variation in deep networks. *arXiv preprint arXiv:1412.6583*, 2014.
- Emily L Denton, Soumith Chintala, Rob Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. In *Advances in Neural Information Processing Systems*, pages 1486–1494, 2015.
- Alexey Dosovitskiy and Thomas Brox. Generating images with perceptual similarity metrics based on deep networks. *arXiv preprint arXiv:1602.02644*, 2016.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.
- J. Green, P. Schechter, C. Baltay, R. Bean, D. Bennett, R. Brown, C. Conselice, M. Donahue, et al. Wide-Field InfraRed Survey Telescope (WFIRST) Final Report. *ArXiv e-prints*, August 2012.

- C. Hirata and U. Seljak. Shear calibration biases in weak-lensing surveys. *Monthly Notices of the Royal Astronomical Society*, 343:459–480, August 2003. doi: 10.1046/j.1365-8711.2003.06683.x.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Diederik P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In *Advances in Neural Information Processing Systems*, pages 3581–3589, 2014.
- Diederik P. Kingma, Tim Salimans, and Max Welling. Improving variational inference with inverse autoregressive flow. *arXiv:1606.04934*, 2016.
- T. Kitching, S. Balan, G. Bernstein, M. Bethge, S. Bridle, F. Courbin, M. Gentile, A. Heavens, et al. Gravitational Lensing Accuracy Testing 2010 (GREAT10) Challenge Handbook. *ArXiv e-prints*, September 2010.
- Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, and Ole Winther. Autoencoding beyond pixels using a learned similarity metric. *arXiv preprint arXiv:1512.09300*, 2015.
- R. Laureijs, J. Amiaux, S. Arduini, J. . Auguères, J. Brinchmann, R. Cole, M. Cropper, C. Dabin, L. Duvet, A. Ealet, and et al. Euclid Definition Study Report. *ArXiv e-prints*, October 2011.
- LSST Science Collaboration, P. A. Abell, J. Allison, S. F. Anderson, J. R. Andrew, J. R. P. Angel, L. Armus, D. Arnett, S. J. Asztalos, T. S. Axelrod, and et al. LSST Science Book, Version 2.0. *ArXiv e-prints*, December 2009.
- R. Mandelbaum, C. M. Hirata, U. Seljak, J. Guzik, N. Padmanabhan, C. Blake, M. R. Blanton, R. Lupton, and J. Brinkmann. Systematic errors in weak lensing: application to SDSS galaxy-galaxy weak lensing. *Monthly Notices of the Royal Astronomical Society*, 361:1287–1322, August 2005. doi: 10.1111/j.1365-2966.2005.09282.x.
- R. Mandelbaum, B. Rowe, J. Bosch, C. Chang, F. Courbin, M. Gill, M. Jarvis, A. Kannawadi, et al. The Third Gravitational Lensing Accuracy Testing (GREAT3) Challenge Handbook. *The Astrophysical Journal Supplement*, 212:5, May 2014. doi: 10.1088/0067-0049/212/1/5.
- Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- B. T. P. Rowe, M. Jarvis, R. Mandelbaum, G. M. Bernstein, J. Bosch, M. Simet, J. E. Meyers, T. Kacprzak, et al. GALSIM: The modular galaxy image simulation toolkit. *Astronomy and Computing*, 10:121–150, April 2015. doi: 10.1016/j.ascom.2015.02.002.
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *arXiv preprint arXiv:1606.03498*, 2016.
- N. Scoville, H. Aussel, M. Brusa, P. Capak, C. M. Carollo, M. Elvis, M. Giavalisco, L. Guzzo, et al. The Cosmic Evolution Survey (COSMOS): Overview. *The Astrophysical Journal Supplement Series*, 172:1–8, September 2007. doi: 10.1086/516585.
- Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems*, pages 3483–3491, 2015.
- Kyle W Willett, Chris J Lintott, Steven P Bamford, Karen L Masters, Brooke D Simmons, Kevin RV Casteels, Edward M Edmondson, Lucy F Fortson, et al. Galaxy zoo 2: detailed morphological classifications for 304 122 galaxies from the sloan digital sky survey. *Monthly Notices of the Royal Astronomical Society*, page stt1458, 2013.
- Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.